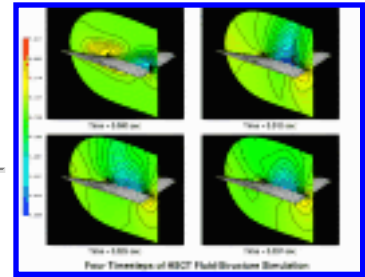
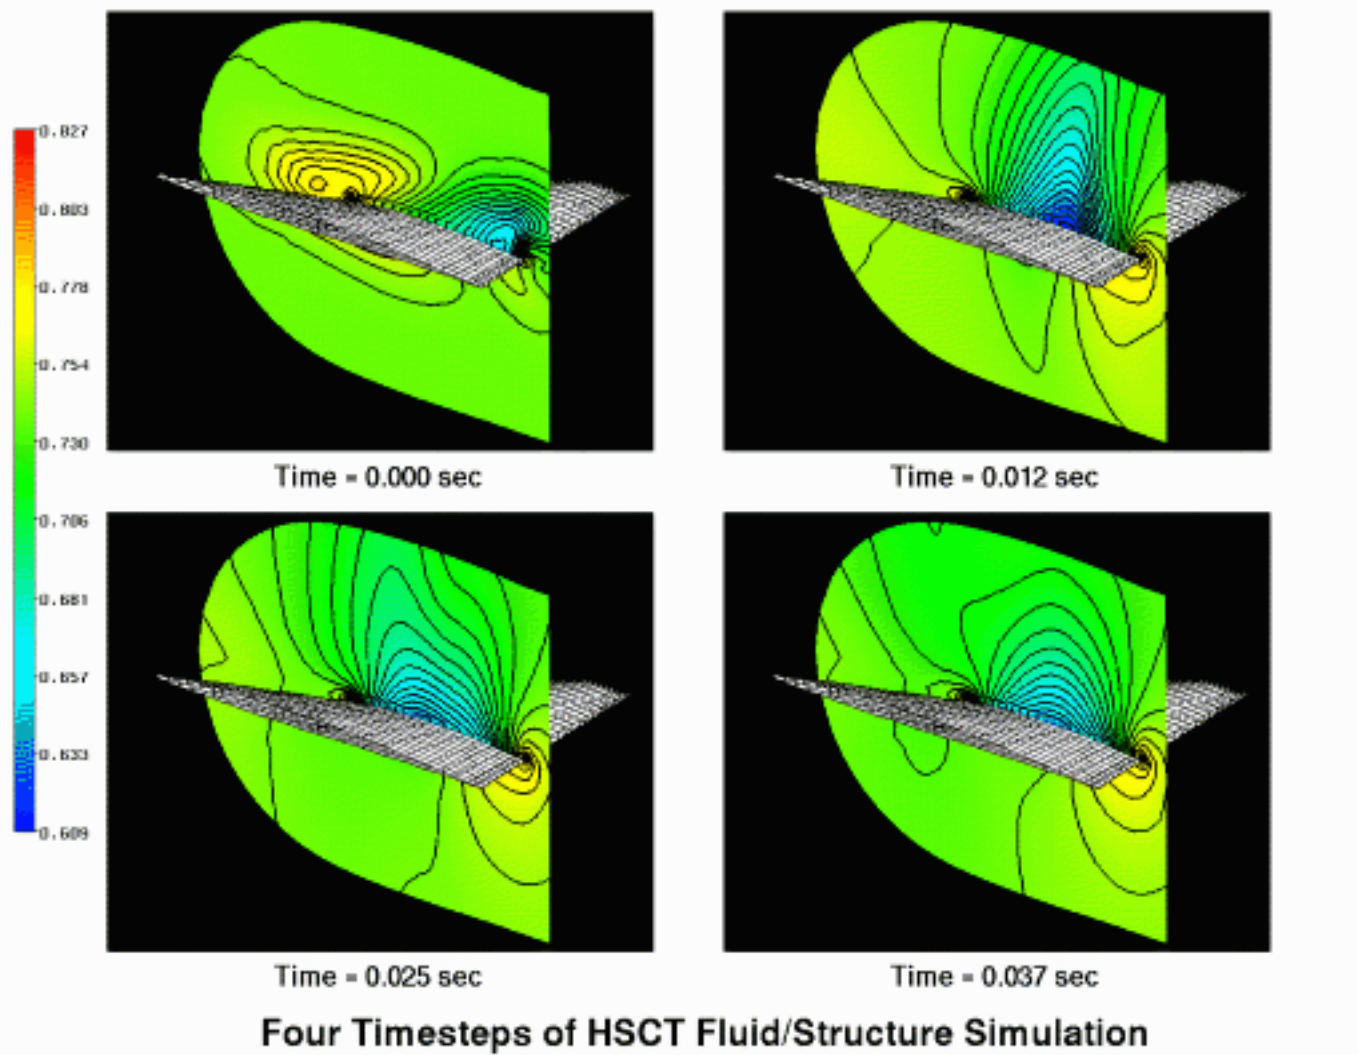
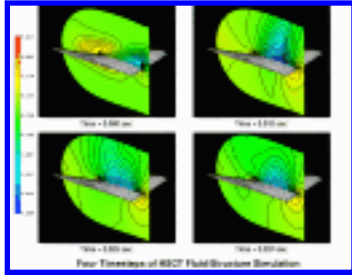


May-June 94 Table of Contents

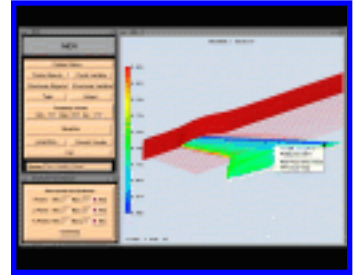
- [Intelligent User Interface is Key to Visualization Tool's Use](#)
- [NAS Links Wind Tunnels to Remote Industry Sites](#)
- [Spectral Methods Improve Grid Partitioning](#)
- [FY94 NASA Research Announcements Awarded](#)
- [Collaboration Tools Help Remote and Local Users](#)
- [Cooperative Research Agreement - A New Approach](#)
- [Parallel UFAT Improves Performance Time](#)
- [Video Teleconferencing Offered](#)
- [Parallel Systems NOP is June 6](#)
- [NAS Saves \\$23K by Publishing User Guide Online](#)
- [Credits For This Issue](#)
- [This Issue's Front Page](#)





[Next Article](#)[Contents](#)[Main Menu](#)[NAS Home](#)

Intelligent User Interface is Key to Visualization Tool's



Use

by [Elisabeth Wechsler](#)

A visualization tool that examines data from multidisciplinary simulations is being developed at NAS by Kristina Miceli, a member of the Applied Research Branch's Visualization and User Interface group. The good news -- scientists don't need to be visualization experts to use it.

MDV, or MultiDisciplinary Visualizer, offers an assistant-based, intelligent user interface to remove the complexity of visualizing multidisciplinary datasets. Although other visualization tools are available, their user interfaces are often so complex that scientists need to become visualization experts.

More Realistic Simulations

MDV was designed for visualizing multidisciplinary simulations. The simulation data used to help prototype MDV was generated by Eddy Pramono and Sisira Weeratunga, of the NAS Applied Research Branch. Their simulation, computed on the Intel iPSC/860, involves an aeroelastic, fluid/structure analysis of the High Speed Civil Transport wing.

"It's a more realistic simulation," Miceli said, "since, in real life, the wing would deform from the pressure of the fluid flow and that, in turn, would affect the fluid pressure. You need to look at both disciplines together in order to do a better design."

MDV inputs datasets generated from this simulation and visualizes the results, depicting the interaction between the two physical disciplines. Scientists can capture several fluids and structures datasets in a series of single time steps and create an animation using each time step as a "frame." By then viewing the animation, they can observe the interaction of the two components dynamically.

Helps with Discipline Interaction

"The MDV tool helps scientists understand the interaction of these two disciplines by visualizing them

simultaneously," said Pramono. "For example, we would like to understand what forces cause the twisting of the wing. MDV has helped to investigate a possible cause of this twisting by visualizing a shock wave that moves back and forth on the wing." In addition to helping understand the physics, MDV has proved useful in visually debugging the multidisciplinary analysis software, as well as in communicating results to the scientists.

Manages Multidisciplinary Data Better

Scientists must work with data from different disciplines -- each with different data structures, formats, and standard visualization techniques. Visualizations from each discipline must be generated individually and then pieced back together for combined analysis, Miceli explained. Two problems with this approach are increased time using visualization tools and a risk of developing poor representations.

The intelligent user interface component of MDV is built on top of a data model, which manages the multidisciplinary data. These datasets can be very large -- on the order of tens of megabytes per time step. In unsteady flow computations, even a small number of time steps may create gigabytes of data. When multiple disciplines are involved in a single simulation, the dataset size can increase by a factor of two or more per discipline, according to Miceli.

Contains System User Information

Visualization and domain knowledge is embedded into the MDV data model, represented as facts and rules. In addition to the data model, a user model contains pertinent information about system users. This information can include the user's background, level of expertise, and any personal preferences (such as preferring a specific color table). Eventually, Miceli plans to adapt the user model to take into account both novice and expert users.

A machine model contains information about the hardware and software characteristics of the user's system that might affect the resulting visualization. Information in the machine model can include the type of system, monitor resolution, graphics capabilities, color facilities, CPU, memory, and available graphics libraries. The mechanism for creating a visualization is an iterative process of refinement that evaluates the relevant information for a data analysis task.

Tasks Affect Choice of Technique

Different tasks performed on the same data by the same user on the same hardware system can require different visualizations. For example, if the scientist is interested in debugging simulation software, the visualization presented would be much different than a visualization generated for publication. A careful analysis of the scientist's work practices is required in order to design an effective visualization tool.

MDV "knows" what types of visualization techniques are appropriate to help scientists do what they

want, Miceli said. For example, if a scientist is interested in viewing the effects of pressure, then a scalar technique -- such as a color contour -- may be appropriate, she explained, adding that other techniques may be more appropriate for other fields.

This knowledge has been built into the tool so that selection of specific techniques is automated. Scientists interested in using MDV can specify:

- the simulation dataset they are interested in viewing
- field variables, such as pressure and deformation
- a subset of the data, such as a slice of the wing
- the desired output type, such as a black and white printout
- the data analysis task -- for example, debugging the code or generating images for a presentation

SuperGlue Supports Rapid Prototyping

MDV is an application code written in SuperGlue, an object-oriented programming environment that supports rapid prototyping of visualization tools. SuperGlue, based on Scheme (a dialect of the language LISP), is being developed by Jeff Hultquist and Eric Raible, both of the NAS Applied Research Branch.

Scheme allows new code to be developed in the background so that a programmer doesn't need to exit SuperGlue to edit, compile, or run the code. This feature allows code to be tested as it is written, Hultquist said. SuperGlue also provides a large object-oriented class hierarchy, which contains the necessary components for generating tools for interactive visualization. Miceli was able to leverage these features to develop MDV more quickly for use with unsteady data.

The main concepts of the object-oriented paradigm include abstraction, encapsulation, and inheritance. Abstraction breaks detailed systems into simple conceptual objects with distinct boundaries. Encapsulation provides a clear separation between the external workings of objects and their internal implementation. Inheritance defines relationships among objects; a class of objects can inherit structure or behavior from another class of objects.

Object-oriented Model Expandable

Because the object-oriented data model is modular, components can be easily added to expand it, Miceli explained. If thermal data is integrated into the data model, for example, an object representing the thermal dataset would be added to the current class structure. Additional methods and attributes can be defined to reflect the specific properties of the new discipline. However, it's still important to gather and

incorporate domain knowledge in order for this tool to be useful, Miceli said.

Another characteristic of the object-oriented data model is its support for metadata, or data about data. Metadata captures information about data format, attributes, and processing history. This capability gives scientists additional high-level information necessary for accurate data analysis, she explained.

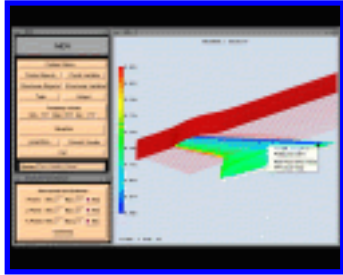
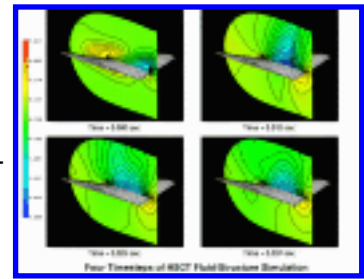
MDV Still a Prototype System

The MDV tool is still a prototype system in the early stages of development. However, it has already shown itself to be useful for the debugging and analysis of multidisciplinary simulations.

For more information about MDV, send email to kmiceli@nas.nasa.gov. For more information about SuperGlue, send email to hultquist@nas.nasa.gov.

[Next Article](#)[Contents](#)[Main Menu](#)[NAS Home](#)

MDV uses standard visualization techniques to view results from unsteady, multidisciplinary simulations. Four time steps of a fluid/structure simulation of the High Speed Civil Transport (HSCT) Wing, computed by Eddy Pramono and Sisira Weeratunga on the Intel iPSC/860, are shown. Fluid pressure data is color-contoured on a fluid grid plane, with wing deformation shown by the sequence of time steps.



This image shows an MDV query to simultaneously visualize pressure on the HSCT wing and velocity of the surrounding fluid flow field. The task chosen is to generate a presentation-style graphic, printed to color hard copy. The result of the query presents the scientist with a visualization of pressure on the surface of the wing via color contouring, and the velocity of the flow field via particle traces and a stream surface. The pop-up menu demonstrates the capability to interactively query the data.



[to the article.](#)

[Next Article](#)[Contents](#)[Main Menu](#)[NAS Home](#)

NAS Links Wind Tunnels to Remote Industry Sites

by [Elisabeth Wechsler](#)

NAS is contributing networking expertise for a collaborative effort between Boeing (Seattle), McDonnell Douglas (Long Beach, CA), and NASA Ames (Moffett Field, CA) that will allow scientists using workstations at remote sites to interact in real time with experiments conducted in Ames' wind tunnels. The concept is to take "the collaborative environment found in research today and bring it into the wind tunnel via local and wide area networks," according to Anthony Lisotta, network development engineer and NAS team lead.

Once the technology has been tested and refined using real computational data and prototype models, Ames will make this technology available "to serve anyone who uses the wind tunnels," said Dave Cooper, Chief, NAS Systems Division.

The impetus for the project came from the Ames Aerodynamics Division, headed by Roy Presley. One of that division's goals has been to increase the productivity of the wind tunnels. Providing remote access not only would reduce travel costs for aerospace industry scientists who need to conduct wind tunnel tests, but would enable an entire research team to view the experiment in progress and make changes simultaneously. Another goal is to make it easier to do real-time comparisons of supercomputer-generated Computational Fluid Dynamics (CFD) data with the experimental results achieved in an operating wind tunnel, using the same model.

The NAS team, including Lisotta, Jim McCabe (network development group lead), and Kelly Russell (network development engineer), has provided network connectivity from the aerospace companies' local area networks (LANs) to AEROnet, the wide area network (WAN) used by NAS, as well as from the wind tunnel LANs to AEROnet.

The project is being done in two phases. The first phase, Integration of Numerical and Experimental Wind Tunnel data (called "I_of_Newt"), began in August 1993 and was completed at year end. This phase, which involved Boeing and Ames only, provided the proof of concept demonstrating that LANs and a WAN could be used to provide a collaborative work environment. Phase 1 also realized Ames' goal for achieving real-time comparisons inside the 11-foot transonic Unitary Plan Wind Tunnel control room, as well as at remote sites, according to Dennis Koga, of the Ames Aerodynamics Division, and principal investigator for Phase 1.

The main challenge of Phase 1 for the NAS team was the short time frame, McCabe said, adding that the NAS team replaced old fiber optic cable and installed new networking devices in the wind tunnel building in about two months.

The second phase, which will add Remote Access Wind Tunnel (RAWT) capability to the I_of_Newt project, began in January and is scheduled to begin tests in the 9-by-7-foot supersonic Unitary Plan Wind Tunnel in mid- June. John Schreiner, also of the Ames Aerodynamics Division, is the RAWT coordinator for Phase 2.

The goals for Phase 2 are to provide real-time video of the model inside the wind tunnel to remote Boeing and McDonnell Douglas sites, allow their scientists to communicate via workstation during the tests, and provide FAST applications across the networks to let researchers compare CFD data with wind tunnel experimental results. The project uses Silicon Graphics Inc. Indigo and Indy workstations, with cameras imbedded to provide real-time visual and voice capability, and InPerson software.

After the completion of Phase 2 this summer, any aerospace company connected to AEROnet can use the remote access wind tunnel capability simply by having the NAS team reconfigure the Wellfleet routers -- a fairly simple task, according to McCabe -- to ensure that each subsequent test suite retains its proprietary nature. For information about requesting remote access wind tunnel technology, send email

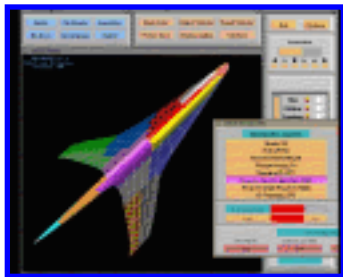
to koga@ames.arc.nasa.gov.

[Next Article](#)

[Contents](#)

[Main Menu](#)

[NAS Home](#)

[Next Article](#)[Contents](#)[Main Menu](#)[NAS Home](#)

Spectral Methods Improve Grid Partitioning for Large- Scale Problems



by Jim Ruppert

Recently developed spectral methods for grid partitioning are speeding up the solution of large-scale problems on the CM-5 and other distributed-memory parallel computers. These spectral algorithms, developed at NAS by Horst Simon, and others, are used to partition a large problem into subproblems for each processor of a parallel computer.

Based on prototype software developed by Simon, Zdenek Johan at Thinking Machines Corp. (TMC) has implemented the algorithm for the CM-5. The software is now available to users as part of the Connection Machine Scientific Software Library (CMSL).

At the Parallel Finite Element Computations Symposium, held October 24-27 in Minneapolis, a number of researchers reported success in using this routine for applications, such as finite element methods for solving Computational Fluid Dynamics (CFD) problems or shallow water equations.

Current research by Simon and Steve Barnard of Cray Research Inc. (CRI) focuses on improving the efficiency and generality of spectral methods.

Grid Partitioning

A computational mesh, such as an unstructured grid, underlies many applications -- for example, numerical fluid dynamics simulations. Time-efficient solutions to very large problems demand a parallel computer, which in turn requires that the computational grid be cut (partitioned) into pieces for each processor to work on. A partition of equal-sized subgrids balances the load by keeping all processors busy.

During the computation, information along the shared subgrid boundaries must be exchanged between processors. Since this communication may be resource costly, the objective is to find a partitioning method that minimizes these boundaries.

Spectral Partitioning Algorithm

Spectral bisection is a technique for cutting a grid into two pieces, such that the nodes are evenly divided and few edges are cut. A good location for the cut is determined by using the connectivity structure of the grid. The resulting pieces can then be cut recursively to partition the grid among a larger number of processors. This is called recursive spectral bisection (RSB).

The effectiveness of the RSB algorithm results from its use of spectral (eigenvector) analysis of the grid. The nodal adjacencies are represented in a matrix (the Laplacian matrix), and the second eigenvector of this matrix is computed. The eigenvector yields an ordering of the nodes from which a cut is determined.

Technical details of the RSB algorithm can be found in "Partitioning Sparse Matrices With Eigenvectors of Graphs," *SIAM Journal of Matrix Analysis and Applications*, Vol. 11 (1990).

Here is a physical interpretation for why the RSB algorithm works: Imagine the grid as a network of springs, which, when tapped with a hammer, will oscillate in various modes (for example, up-down or left-half up, right-half down). An eigenvector corresponds to one of these modes, and gives a displacement value for each node on the grid. The simplest of these modes corresponds to the second eigenvector, which the RSB algorithm uses to partition the nodes into two groups -- those above and those below the median displacement -- a natural division that tends to minimize the number of edges cut.

The RSB algorithm was compared favorably with two earlier methods, recursive coordinate bisection and recursive graph bisection, in "Partitioning of Unstructured Problems for Parallel Processing," *Computing Systems in Engineering*, Vol. 2 (1991). The performance measure was the number of edges cut, since this indicates the interprocessor communication cost. On two-dimensional and three-dimensional problems of up to 45,000 vertices, RSB was found to produce a smaller cut than either of the other algorithms, and was twice as effective in some cases.

Implementation

The RSB algorithm, as implemented by Johan on the CM-5 machine, is a parallelized CMSSL routine. On three-dimensional grids with 80,000-100,000 vertices, a 128-way partition can be computed in less than a minute on a 32-node CM-5. For more information about this implementation, see "An Efficient Communication Strategy for Finite Element Methods on the Connection Machine CM-5 System," TMC *Technical Report TR-256* (1993).

Generally, the basic RSB algorithm has a long execution time when used for grids larger than 100,000 vertices. Barnard, of CRI, and Simon developed a sequential implementation of RSB that relies heavily on multigrid algorithms and runs up to 20 times faster than the basic RSB code. On a standard SGI workstation, grids with 30,000-40,000 vertices can be partitioned in one or two minutes. For more

information about this method, see "A Fast Multilevel Implementation of Recursive Spectral Bisection for Partitioning Unstructured Problems," [NAS Technical Report RNR-92-033](#).

RSB Available in Tools

RSB is also available as an option in TOP/DOMDEC, a tool developed by Charbel Farhat of the University of Colorado at Boulder, which offers interactive software for grid partitioning and parallel processing. In addition to RSB, it offers several other partitioning algorithms.

TOP/DOMDEC can also improve the quality of a grid partition using several optimization algorithms. The user interface includes three- dimensional graphics, an interprocessor communication simulator for today's massively parallel systems, and an output function with parallel I/O data structures. The basic features of TOP/DOMDEC and their application to the parallel solution of computational fluid and solid mechanics problems are described in "TOP/DOMDEC -- A Software Tool for Mesh Partitioning and Parallel Processing," [NAS Technical Report RNR-93-011](#).

The RSB algorithm is also available as part of the Chaco software package for grid partitioning, developed at Sandia National Laboratories by Bruce Hendrickson and Robert Leland. For detailed information about this package, see *The Chaco User's Guide -- Version 1.0*.

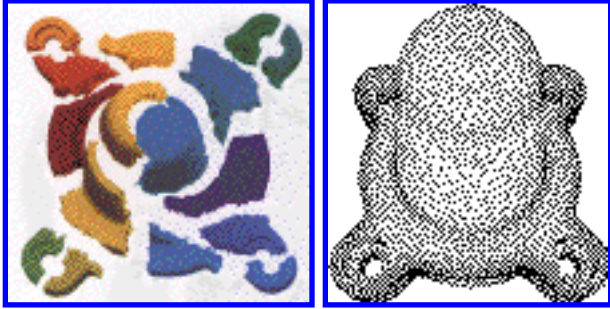
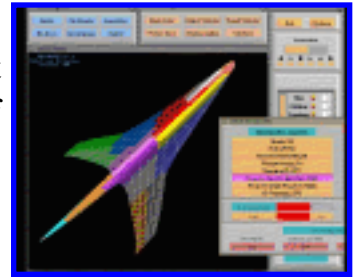
Future Work

Research in progress on spectral methods includes an investigation of a fast update mechanism for dynamically changing grids, as well as theoretical studies for optimizing RSB.

The reports mentioned in this article are available from the [NAS Documentation Center](#). For other questions, send email to: **simon@nas.nasa.gov**.

[Next Article](#)[Contents](#)[Main Menu](#)[NAS Home](#)

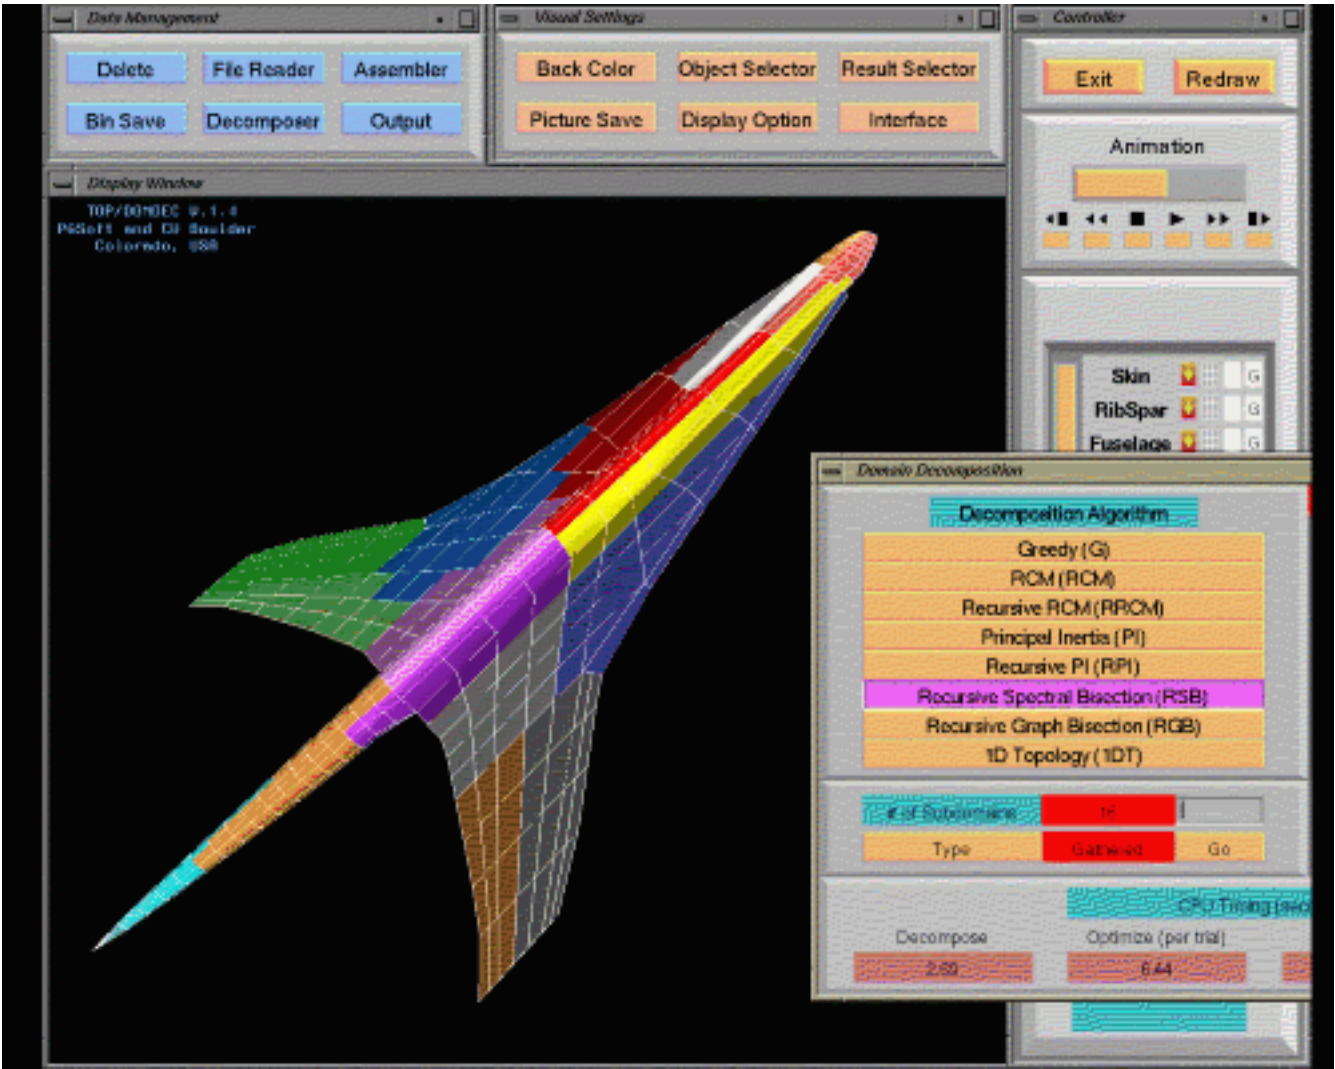
TOP/DOMDEC interface, showing 16-way partition of High Speed Civil Transport (HSCT) structural analysis grid produced using the high-density grid at the nose. The HSCT model was developed by Charbel Farhat of the University of Colorado.



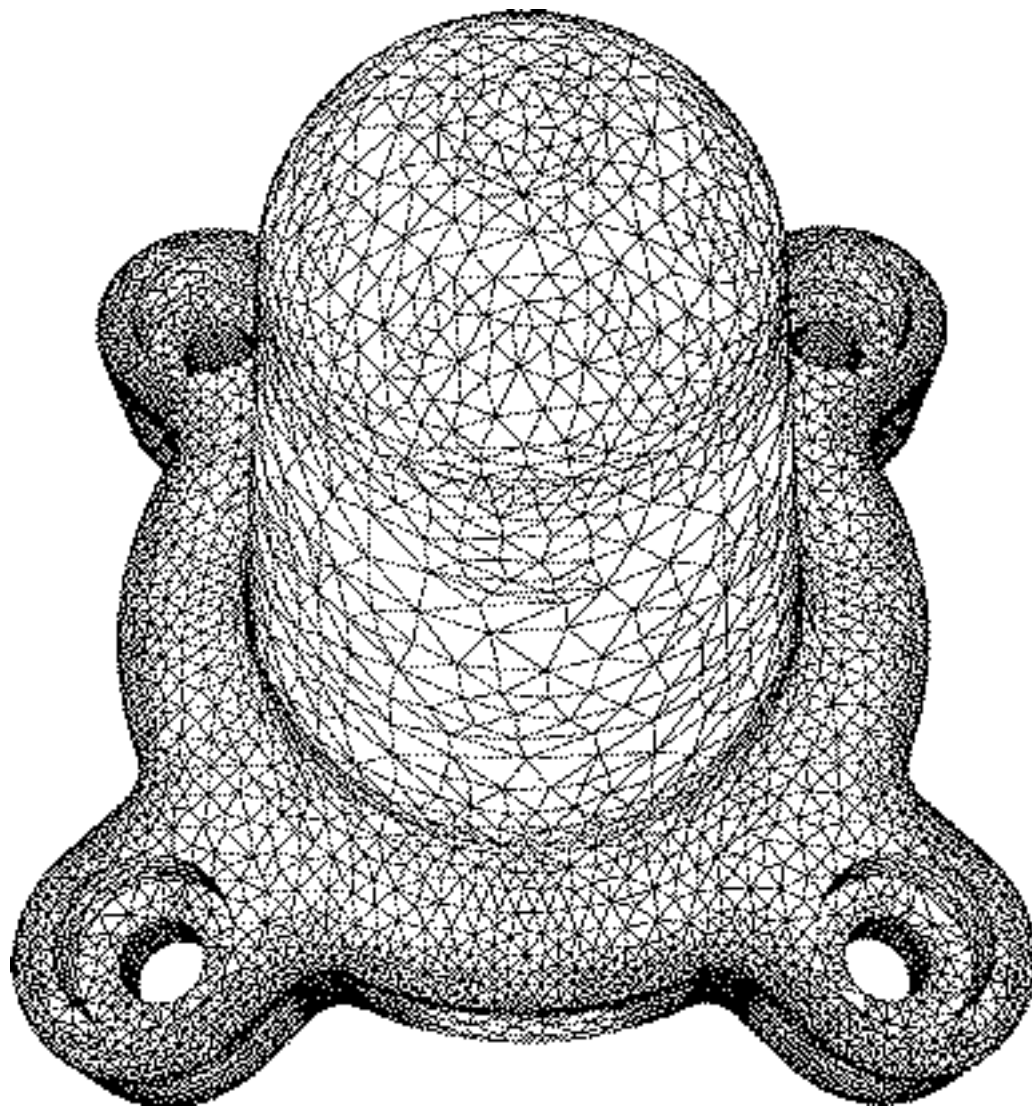
Two views of an assembly part, showing unstructured grid of tetrahedra, and 16-way grid partition produced by the Connection Machine Scientific Software Library implementation of recursive spectral bisection algorithm. Several cuts pass through object holes to reduce the number of cut edges. The grid, consisting of 81,649 tetrahedra, was produced with the octree grid generator developed by Mark Shephard and others at Rensselaer Polytechnic Institute.



[to the article.](#)







[Next Article](#)[Contents](#)[Main Menu](#)[NAS Home](#)

FY94 NASA Research Announcements Awarded

by [Jill Dunbar](#)

Two FY94 NASA Research Announcements (NRA) have resulted in awards to ten institutions for research in support of the NAS Program's objectives in Computational Aerosciences.

The two broad areas of research covered by the NRA are parallel algorithms/systems software and scientific visualization. Of this year's awards, four are for new projects and six are for continuation projects funded through last year's NRA. The four new awards were split evenly between the parallel algorithms and visualization categories. A total of 42 proposals were received and evaluated by scientists in the NAS Applied Research Branch.

Parallel Projects Focus on CFD

Funding for two new proposals in the parallel algorithms area was awarded to the University of Illinois and the University of Minnesota.

The goal of principal investigator Youcef Saad at the University of Minnesota Computer Science Department is to investigate the applicability of some recent developments in numerical linear algebra to the solution of large sparse systems of linear equations arising in the simulation of three-dimensional flow fields associated with aerospace vehicles. Saad plans to investigate the applicability of a set of preconditioners for the solution of linear systems associated with realistic, 3D Computational Fluid Dynamics (CFD) simulations. The successful completion of this investigation would lead to significant improvements in the efficiency and reliability of CFD simulations on high-speed computers.

Laxmikant Kale, principal investigator at the University of Illinois Department of Computer Science, has proposed investigating whether object-oriented programming methodology can help control the complexity of parallel programming in CFD. This work involves implementing the NAS Parallel Benchmarks suite using Charm ++, a portable concurrent object-oriented programming language (based on C++) developed at Illinois. For each benchmark, the investigator will evaluate performance results obtained on a variety of parallel and distributed architectures. The programming process involved in the implementation and optimization of the benchmarks will also be evaluated to understand the effectiveness of the programming methodology.

"Viz" Projects Take Fresh Approach

Two new visualization projects, awarded to Stanford University and San Jose State University, take a fresh approach to visualizing tensor fields encountered in continuum mechanics.

The work at Stanford seeks a "unified framework for the visualization of vector and tensor fields." It will build on previous NASA-supported vector and tensor field visualization research conducted by principal investigator Lambertus Hesselink, according to Arsi Vaziri, one of the key NAS participants in the NRA process. Hesselink and student researchers plan to create elementary surface tensor icons and improve existing global tensor icons, then extend these results to time-dependent datasets.

For surface tensor icons, Stanford will study techniques for representing 3x3 tensor data on two-dimensional surfaces. Using textures and topological approaches, researchers hope to improve global tensor icons, based on their previous work on topological decomposition of vector datasets. New techniques for depicting and analyzing the topology of 2D tensor fields will be combined with renderings of vector fields.

R.K. Dodd, principal investigator at San Jose State's Department of Mathematics and Computer Science, will focus on an alternative method of visualizing symmetric tensors arising from the study of fluid flow. His work will use a Riemannian metric derived from the tensor field. Within this framework, Dodd plans to calculate and visualize natural geometric invariants (such as curvature) of the metric by using an orthonormal basis defined over the space, in the manner of the late French geometer Elie Cartan.

Continuing Grants

Six grants were awarded to continue work begun last year in support of the NAS Computational Aerosciences research program. Awards went to: Charbel Farhat, University of Colorado; Mark Shephard, Rensselaer Polytechnic Institute; and V. S. Sunderam, Emory University, for research in parallel algorithms and system software. For research in scientific visualization, continuation grants were awarded to: Deborah Silver, Rutgers University; Michael Stokes, Mississippi State University; and Andries van Dam, Brown University. A continuing cooperative research grant with Jane Wilhelms at the University of California at Santa Cruz was funded for an additional year.

A NASA Research Announcement for FY95 is expected to be issued later this year. Look for information in a future issue of *NAS News* or contact Arsi Vaziri at vaziri@nas.nasa.gov.

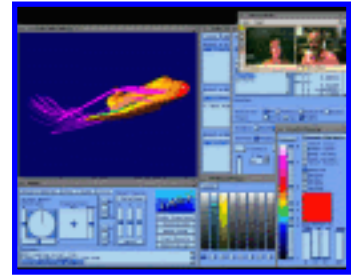
[Arsi Vaziri](#) contributed to this article.

[Next Article](#)[Contents](#)[Main Menu](#)[NAS Home](#)

[Next Article](#)[Contents](#)[Main Menu](#)[NAS Home](#)

Collaboration Tools Help Remote and Local Users

by [Jill Dunbar](#)



Recently, a combination of two tools were tested for improving communications between scientists at different locations and between managers at Ames Research Center and NASA headquarters in Washington, D.C. The tools that were tested are InPerson, a desktop video conferencing tool from Silicon Graphics Inc. and FAST, the Flow Analysis Software Toolkit, a scientific visualization tool developed at NAS.

The tests were conceived by Val Watson, Ames Fluid Dynamics Division, to illustrate the potential of these tools to improve communication and collaboration, and to measure the potential impact that these tools have on networks. Watson, who works closely with NAS technical staff on FAST and other visualization techniques, has long been an advocate of using workstations for collaboration and instruction. That, coupled with his concern about clogged networks and the corresponding reduction in the users' ability to get information, led him to propose testing this combination of tools.

The tests indicate that remote tools can effectively eliminate the "barrier of distance" for many types of collaboration and can also improve local collaborative work. For some types of work, these tools can be even more effective than in-person meetings, because each participant has a work-station at their disposal during a collaborative session. Plus, at the end of the session, each person has the results of the collaboration on his or her own workstation in a form that can be conveniently used or modified.

Tools Provide Eye-to-Eye Contact

During the tests, InPerson provided small live video pictures of participants for eye-to-eye contact, or a small picture of any object from remote sites. FAST provided the capability to jointly analyze high-resolution, dynamic, three-dimensional data. Examples of data analyzed during the tests include a flow field around the Space Shuttle (from Ames), atmospheric pollution over the eastern half of the United States (from the Environmental Protection Agency), and a vortex breakdown (from the University of California at Davis).

The two tools were linked via Internet for data communications. NSI links between Ames and NASA headquarters were used during the test period. At least one of these links was limited to T1 bandwidth, which transmits at 1.5 megabits per second (mb/sec).

Test Results Highlight Efficiency

Tests results show that although the small video pictures provided by InPerson are blotchy when participants move quickly, the pictures are good when participants stay relatively still -- as shown in the accompanying figure. Picture quality is entirely adequate to get the sense of eye-to-eye contact and to read participants' expressions. This ability to read another person's acceptance or rejection of information is a major factor in making the collaboration highly effective. Participants agreed that being able to see the other person added significantly to the quality of communication.

The remote visual analysis using FAST is also very effective -- in fact, remote analysis proved to be nearly the same as if all participants were together looking at the same workstation. Additionally, using FAST in a remote session is almost identical to using it in stand-alone mode, so retraining for users is not required.

InPerson Impacts Network

Using the InPerson desktop video conferencing tool has a substantial impact on the network. Test measurements show that the average bandwidth used was .25 mb/sec. and the peak bandwidth was .50 mb/sec. On average, InPerson used one-sixth of the total bandwidth of the T1 links between Ames and NASA headquarters. However, it was found that a much lower average bandwidth could be used by "freezing" the video whenever eye-to-eye contact was not necessary. For example, the video can be frozen during periods when participants are concentrating on the visual analysis of 3D data.

Other information gathered from the tests show that using FAST has a negligible impact on the network -- only one kilobit per second of bandwidth was used for remote visualization. FAST does not send information as pixels over the network. Instead, FAST transmits only the user instructions on how to visualize the data, and a copy of FAST on each workstation generates the pixels locally for each participant. The raw data to be analyzed is sent before the visual analysis starts. In the tests using Space Shuttle data (7 megabytes), about three minutes were required to send the data prior to visual analysis.

Lessons Learned

NAS researchers learned valuable information about remote and local collaboration combining InPerson and FAST.

- The tools are highly effective. They not only can remove the "distance barrier" for remote collaboration, but can improve local collaboration as well.
- Networks will clog if transmission of pixels is not limited to static pictures, short video clips, or small (or low resolution) live video pictures.

Unfortunately, many tools being released for remote collaboration rely on sending most information as pixels. Sending pixel information to generate the dynamic pictures that users see on workstation monitors would require 2 gigabits per second -- twice the bandwidth of the proposed backbone of the "information super highway." Even if pixel data were compressed by a factor of 1000, the bandwidth would be too high to send over typical T1 lines. It would take only a few people using tools that rely on sending most information as pixels to clog networks.

- High resolution, dynamic, 3D scenes can be used in collaborative sessions without using much bandwidth by using tools such as FAST, which sends the information in a more efficient form than pixels.
- The most effective collaboration tools will be based on joint control of tools that are currently used for scientific work, rather than "shared whiteboard" type tools, which allow participants to jointly use paint and draw type tools to create images. With FAST, scientists can do remote joint visual analysis and conduct further independent visual analysis, modification, or animation.

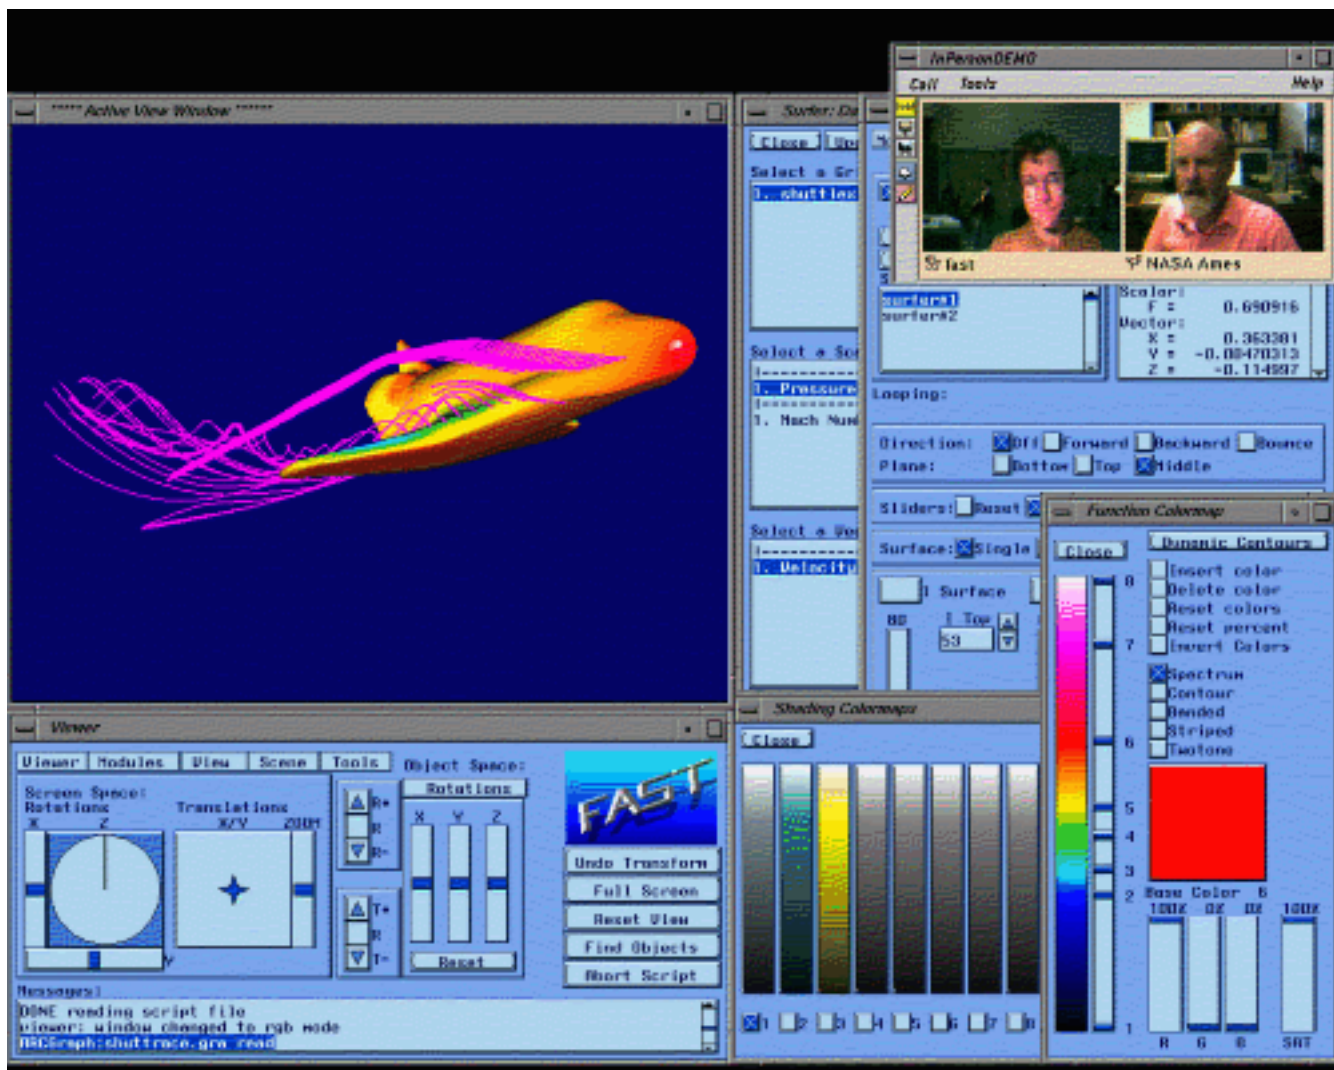
For More Information

More testing with three or more concurrent sites is being planned. To find out how to use FAST remotely, contact John West at jwest@nas.nasa.gov. For more information on using these collaboration tools, contact Val Watson at watson@nas.nasa.gov.

[Val Watson](#) and [Pamela Walatka](#) contributed to this article.

Staff members in the NAS Systems Development Branch provided on-site support for the collaborative tests: Jude George provided network measurements; Kevin McCabe and John West developed most of the utilities for remote control of FAST. Ron Bailey (Director, Ames Office of Aerophysics), Pam Richardson and Robert Whitehead (both of the Office of Aeronautics, NASA headquarters), and Paul Kutler (Chief, Ames Aerodynamics Division) donated time to test the tools. Mike Little and Judy Pellerin from NASA headquarters also provided major support. Silicon Graphics provided support for InPerson.

[Next Article](#)[Contents](#)[Main Menu](#)[NAS Home](#)



A view of the FAST screen, with various functions and options for providing remote collaborative visual analysis to scientists. Collaboration is improved with desktop video conferencing provided by InPerson software (upper right in figure). The blurred image on the left, captured at a NAS demonstration of remote collaboration at Supercomputing '93, shows the quality of transmission of a moving subject. At the other end of the connection, the image of Val Watson (upper right) was "frozen" to reduce network bandwidth.



[to the article.](http://www.nas.nasa.gov/Pubs/NASnews/94/05/val_light.html)

[Next Article](#)[Contents](#)[Main Menu](#)[NAS Home](#)

Cooperative Research Agreement - A New Approach

by [Elisabeth Wechsler](#)

The NAS Systems Division is among the first NASA organizations to implement the Cooperative Research Agreement (CRA), a vehicle developed by NASA headquarters that combines some aspects of the NASA Research Announcement (NRA) with some aspects of a government procurement.

The goal behind this CRA is "to advance U.S. leadership and competitiveness in high-performance computing in general and in computational aerosciences in particular," according to David H. Bailey of the NAS Applied Research Branch. Up to \$25 million in funding is being provided by the High Performance Computing and Communications Program.

The CRA is a proposal process to solicit groups of research entities (for example, hardware vendors and their personnel, researchers, universities, aerospace companies, government laboratories, and independent software vendors) to put together a package of research and equipment.

Equipment 'Loaned,' Not Purchased

One important feature of the CRA, however, is that "we're not taking ownership of the equipment that's provided," Bailey said. "It is understood that the vendor will take away the equipment at the end of the agreement."

The invitation to submit proposals for consideration under the CRA went out on July 1, 1993, and a preproposal conference was held at NAS on August 3 to explain the submission procedures. The deadline for submissions to NAS was December 1. The CRA stipulates no written constraints for the project's duration, but Bailey said that the research and development effort is expected to last at least two years.

Research Key to Selection Process

The proposals include specific research projects to be performed by each submitting group on particular equipment. "The evaluation of those research proposals is an important part of the overall proposal evaluation," Bailey said, emphasizing that "this is not just a hardware solicitation."

A technical evaluation team, headed by Bailey, has been reviewing the responses since December, and an

announcement of the winning team or teams is expected soon. The selected team(s) will work with NASA as an active partner, and together will form a consortium. NASA will monitor the research and development being done, and it is expected that staff from all NAS branches will be involved in making the computer equipment more usable. This may include work on operating systems, utilities, languages, and tools, Bailey said.

Each Party Will Contribute Resources

The CRA is based on the principle that "each of the different parties will put up some resources -- manpower, funds, or equipment -- toward the combined effort," Bailey said. For example, NAS will provide funds and research time but will receive no direct funding. Other organizations will do some other work "beyond what they're paid for," he added.

[Next Article](#)[Contents](#)[Main Menu](#)[NAS Home](#)

[Next Article](#)[Contents](#)[Main Menu](#)[NAS Home](#)

Parallel UFAT Improves Performance Time

A parallel version of the Unsteady Flow Analysis Toolkit (UFAT) completed development in April and is now being tested. UFAT is a three-dimensional particle tracing software tool for large-scale, time-dependent, Computational Fluid Dynamics (CFD) flow data. Version 2.2 was developed to take advantage of the parallel processing capabilities available at NAS.

Tests show that performance can be improved by several factors -- depending on the number of processors used -- according to David Lane, who developed both the sequential and parallel versions of UFAT. For example, a single processor on the Convex C3240 (**vonkarman**), took 68 minutes of wall-clock time to compute particles from 200 time steps of a Delta Wing flow dataset, which consisted of 900,000 grid points. When four processors were used, the time was reduced to 27 minutes.

The parallel version will also be available on the CRAY C90 (**vonneumann**) and the four support processing systems **amelia**, **fred**, **orville**, and **wilbur**.

Factors Affecting Performance

The increase in UFAT's performance depends on the following factors:

- the number of time steps used
- the disk I/O rate to read the grid and solution file
- the number of seed points (the locations from which particles are released)

Since particles are traced through all time steps, the number of time steps determines how long UFAT will execute. At each time step, UFAT reads the solution file as well as the grid file, if the grid is unsteady. To read a grid with several million grid points can take several seconds to several minutes, depending on the system's disk I/O performance rate. A set of new particles are released from the user-specified seed locations at each time step, so the number of particles increases linearly at each new time step. As the number of particles increases, so does the time to trace the current particles.

Particle Tracing Easily Parallelized

Particle tracing is an embarrassingly parallel application -- that is, it is easy to "parallelize" the code without extensive rework -- because the traces can easily be computed independently of each other.

Scientists in the Ames Fluid Dynamics Division use UFAT to compute streaklines from their unsteady flow fields. Among them are Chris Atwood, who has performed computations of a controlled store separation from a cavity; Goetz Klopfer and Shigeru Obayashi, who have evaluated a new zoning method for grids with time-varying boundary conditions; and Robert Meakin, whose development of overset grid methods for the V22 Tiltrotor aircraft was featured in an article on UFAT in the July-August 1993 issue of *NAS News*.

What's Next?

A distributed Parallel Virtual Machine-based version of UFAT is planned for release at the end of this year. This version will use a message-passing library to communicate data between systems, and will provide a graphical user interface.

For more information about the parallel version of UFAT, send email to lane@nas.nasa.gov.

[Next Article](#)[Contents](#)[Main Menu](#)[NAS Home](#)

[Next Article](#)[Contents](#)[Main Menu](#)[NAS Home](#)

Video Teleconferencing Offered

NAS recently made multiparty video teleconferencing available to its commercial aerospace customers -- at no charge -- for participating in the NAS User Seminars and the NAS User Group (NUG) meetings, both held for one hour each month. Customers may participate on a per-event basis, depending on each site's interest in a particular topic.

So far, three commercial sites -- Lockheed (Fort Worth), McDonnell Douglas (St. Louis), and Rockwell International (Thousand Oaks, CA) -- have used their existing technology to take advantage of this capability. Video tapes created during broadcast time at each site are available to employees who can't attend the actual session.

"Overall, the people who attended the [user] seminars really liked it," said Ray Cosner, manager for CFD applications at McDonnell Douglas. He added that the NUG meetings are "good and useful in terms of keeping in touch with what's going on."

Video teleconferencing facilities at commercial sites must have the capability of connecting to the U.S. Sprint network. Eight new network connections are available for commercial sites in addition to six pre-existing connections to other NASA sites, where video teleconferencing for these events has been available for several years. NASA's Imaging Technology Branch coordinates the technical aspects of the video teleconferences and adjusts the transmission quality.

User Seminars Cover Diverse Topics

On average about 50 participants, including on site and remote viewers, attend the user seminars, according to Marcia Redmond, NAS training coordinator. Seminar dates are available to NAS users directly through *nasgopher* (look in NAS/CCF Training under "course_desc"). Contact Redmond directly at (415) 604-4373 or send email to redmond@nas.nasa.gov.

Most topics are announced about 30 days before the predetermined dates. Seminars cover a diverse range of topics, including visualization tools, Fortran techniques, parallel algorithms and programming, and network technology. Among the seminars given so far this year are "Efficient Parallelization of a Class of Implicit Numerical Algorithms" by Rob Van der Wijngaart and "Technologies for Visualization in Computational Aeroscience" by Kristina Miceli (*see related article on page 1*).

Users Give Feedback Through NUG Meetings

The [NAS User Group](#) meetings are usually held on the second Tuesday of each month starting at 11:00 a.m., Pacific time. The meetings are an opportunity for NAS customers to exchange information with the on-site technical staff about problems and issues, as well as to give input to and keep informed about changes to NAS systems and policies. Recent topics have included the New Operational Period for the high-speed processors, the new Cray Accounting System, network upgrades, and connectivity issues.

For more information about how to participate in NAS-sponsored video teleconferences, contact Rita Williams, NAS Computational Services Branch, at (415) 604-6682 or send email to rita@nas.nasa.gov.

[Next Article](#)[Contents](#)[Main Menu](#)[NAS Home](#)

[Next Article](#)[Contents](#)[Main Menu](#)[NAS Home](#)

Parallel Systems NOP is June 6

The NAS Parallel Systems NOP, or New Operational Period, begins on Monday, June 6, and runs through Sunday, June 4, 1995.

Training classes for the Intel iPSC/860 and Thinking Machines Corp.'s Connection Machine CM-5 parallel systems will be held at NAS in two half- day sessions each during mid-June.

The classes will be videotaped, and copies of the tapes can be obtained by sending an email request to doc-center@nas.nasa.gov. Instructors will be members of the Parallel Systems consulting group, including Ed Hook, Chuck Niggley, Subhash Saini, and Bill Saphir, all of the NAS Computational Services Branch.

The Announcement of Opportunity inviting proposal submissions for the NAS parallel systems was mailed out in March, and proposals have been received and reviewed. All necessary paperwork will be sent to users whose proposals have been approved. June 1 is the deadline for users to submit the Account Request Form, in order to allow time for NAS to prepare accounts for the NOP. Users who have questions about the status of their proposals may contact Pat Elson, User Interface Manager, at (415) 604-4463, or send email to pelson@nas.nasa.gov.

More detailed information about this year's NOP, such as the dates and course outlines for training classes, will be posted on *nasgopher* (look under Parallel Systems for "NOP Information") as it becomes available.

This year, the *NAS User Guide* will be available online only, unless users specifically request a hard-copy version from the NAS Documentation Center. (See "[NAS Saves 23K by Publishing User Guide Online](#)".) A copy of the NAS Processing System Network *NPSN Quick Reference Guide* will be mailed to users who are allocated time on the parallel systems.

[Next Article](#)[Contents](#)[Main Menu](#)[NAS Home](#)

NAS Saves \$23K by Publishing User Guide Online

Putting the 900-page *NAS User Guide* online for the 1994-95 New Operational Period (NOP) saved NAS \$23,632 -- about \$16.88 per manual for approximately 1,400 NAS users -- in production and mailing costs, according to Santa Huguez, the NAS Documentation Librarian.

A six-page *NPSN [NAS Processing System Network] Quick Reference Guide*, known informally as the NAS "hitchhiker's guide," was mailed to users before the NAS High Speed Processor NOP, which began on March 7. This document describes the organization of the *NAS User Guide* and gives tips on where to look for various types of information.

Another benefit of having an online user guide is that changes to technical information can be updated more quickly and easily, and those changes can be made available sooner to users.

The *NAS User Guide* is available through *nasgopher* in either PostScript or FrameMaker formats, allowing users to print the manual in published quality format, if desired.

Last year, NAS saved about \$36,000 by putting Cray documentation online instead of sending copies to the 400 principal investigators who use NAS supercomputers. NAS will continue this practice.

[Next Article](#)[Contents](#)[Main Menu](#)[NAS Home](#)

May-June 1994, Vol. 2, No. 3

Executive Editor: Marisa Chancellor

Editor: Jill Dunbar

Senior Writer: Elisabeth Wechsler

Contributing Writers: David Lane, Jim Ruppert, Val Watson

Other Contributors: David Bailey, Dave Cooper, Jim Crow, Pat Elson, Santa Huguez, Jeff Hultquist, Chris Kleiber, Tom Lasinski, Anthony Lisotta, Jim McCabe, Kristina Miceli, Eddy Pramono, Marcia Redmond, Horst Simon, Sam Uselton, Arsi Vaziri, Sisira Weeratunga, Tom Woodrow

Editorial Board: Marisa Chancellor, Jill Dunbar, Chris Gong, Serge Polevitzky, Jim Ruppert, Pamela Walatka, Elisabeth Wechsler, Rita Williams



NEWS

Volume 2, Number 3

May - June 1994

NAS Links Wind Tunnels to Remote Industry Sites

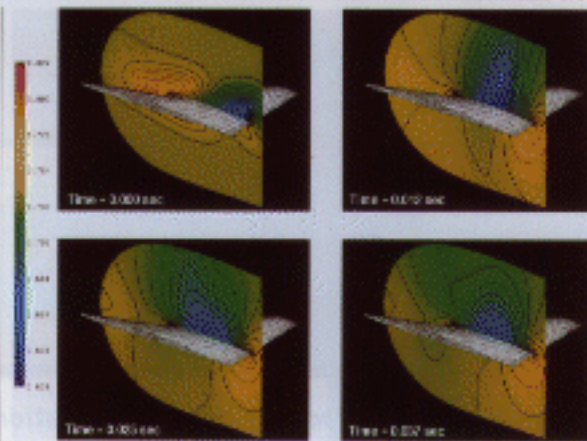
by Elisabeth Wechsler

NAS is contributing networking expertise for a collaborative effort between Boeing, Douglas, McDonnell Douglas, Boeing, Boeing, CA, and NASA Ames-Danell Ford, CA, that will allow scientists using workstations at remote sites to interact in real time with experiments conducted in Ames' wind tunnels. The concept is to take "the collaborative environment found in research today and bring it into the wind tunnel via local and wide area networks," according to Anthony Liotta, network development engineer and NAS team lead.

Once the technology has been tested and refined using real computational data and prototype models, Ames will make this technology available "to serve anyone who uses the wind tunnels," said Dave Cooper, Chief, NAS Systems Division.

The impetus for the project came from the Ames Aerodynamics Division, headed by Ray Zeeley. One of that division's goals has been to increase the productivity of the wind tunnels. Providing remote access not only would reduce travel costs for aerospace industry scientists who need to conduct wind tunnel tests, but would enable an entire research team to view the experiment in progress and make changes simultaneously. Another goal is to make it easier to do real-time comparisons of supercomputer-generated Computational Fluid Dynamics (CFD) data with the experimental results achieved in an operating wind tunnel, using the same model.

Continued on page 7



MDV uses standard resolution techniques to view results from unsteady, multidisciplinary simulations. Four side-angle views of a fluid structure simulation of the High-Speed Civil Transport (HSCT) Wing, generated by Zeeley (Ames) and Dave Zeeley (on the left) (PSCB), are shown. Fluid pressure data is color-coded on a grid plane, with wing deformation shown by the sequence of the images.

Intelligent User Interface is Key to Visualization Tool's Use

by Elisabeth Wechsler

A visualization tool that examines data from multidisciplinary simulation is being developed at NAS by Systems Model, a member of the Applied Research Branch's Visualization and User Interface group. The good news—scientists don't need to be visualization experts to use it.

MDV, or Multidisciplinary Visualizer, offers an intuitive, intelligent user interface to remove the complexity of visualizing multidisciplinary datasets. Although other visualization tools are available, their user interfaces are often so complex that scientists need to become visualization experts.

More Realistic Simulations

MDV was designed for visualizing multidisciplinary simulations. The simulation data used to help prototype MDV was generated by Larry Pramono and Solis Hwangwang, of the NAS Applied Research Branch. Their simulation, compared on the left (PSCB), involves an aerodynamic fluid structure analysis of the High-Speed Civil Transport wing.

"It's a more realistic simulation," Meel said. "In real life, the wing would deform from the pressure of the fluid flow and that, in turn, would affect the fluid pressure. You need to look at both disciplines together in order to do a better design."

MDV inputs datasets generated from this simulation and visualizes the results, depicting the interaction between the two physical disciplines. Scientists can capture several fields

and structure datasets in a series of single time steps and create an animation using such time steps as a "frame." By then viewing the animation, they can observe the interaction of the two components dynamically.

Helps with Discipline Interaction

"The MDV tool helps scientists understand the interaction of these two disciplines by visualizing them simultaneously," said Pramono. "For example, we would like to understand what factors cause the twisting of the wing. MDV has helped to investigate a possible cause of this twisting by visualizing a shock wave that moves back and forth on the wing." In addition to helping understand the physics, MDV has proved useful in visually debugging the multidisciplinary analysis software, as well as in communicating results to the scientists.

Manages Multidisciplinary Data Better

Scientists must work with data from different disciplines—each with different data structures, formats, and standard visualization techniques. Visualizations from such disciplines must be generated individually and then pieced back together for combined analysis. Pramono explained. Two problems with this approach are increased time using visualization tools and a risk of developing poor representations.

The intelligent user interface component of MDV is built on top of a data model, which manages the multidisciplinary data. These datasets can be very large—on the order of

Continued on page 7

THIS ISSUE

Parallel NOP Starts
page 3

Spectral Methods Improve Partitioning
page 4

Video Teleconferencing
page 6